Synergistic Sorting, MultiSelection and Deferred Data Structures on MultiSets

Jérémy Barbay[1], Carlos Ochoa[1], and Srinivasa Rao Satti[2]

[1] Departamento de Ciencias de la Computación, Universidad de Chile, Chile
jeremy@barbay.cl, cochoa@dcc.uchile.cl
[2] Department of Computer Science and Engineering, Seoul National University, South Korea
ssrao@cse.snu.ac.kr

arXiv:1608.06666v2 [cs.DS] 30 Sep 2016

Abstract. Karp et al. (1988) described Deferred Data Structures for Multisets as "lazy" data structures which partially sort data to support online rank and select queries, with the minimum amount of work in the worst case over instances of size n and number of queries q fixed (i.e., the query size). Barbay et al. (2016) refined this approach to take advantage of the gaps between the positions hit by the queries (i.e., the structure in the queries). We develop new techniques in order to further refine this approach and to take advantage all at once of the structure (i.e., the multiplicities of the elements), the local order (i.e., the number and sizes of runs) and the global order (i.e., the number and positions of existing pivots) in the input; and of the structure and order in the sequence of queries. Our main result is a synergistic deferred data structure which performs much better on large classes of instances, while performing always asymptotically as good as previous solutions. As intermediate results, we describe two new synergistic sorting algorithms, which take advantage of the structure and order (local and global) in the input, improving upon previous results which take advantage only of the structure (Munro and Spira 1979) or of the local order (Takaoka 1997) in the input; and one new multiselection algorithm which takes advantage of not only the order and structure in the input, but also of the structure in the queries. We described two compressed data structures to represent a multiset taking advantage of both the local order and structure, while supporting the operators `rank` and `select` on the multiset.

Keywords: Deferred Data Structure, Divide and Conquer, Fine Grained Analysis, Quick Select, Quick Sort, Compressed Data Structure, Rank, Select.

1 Introduction

Consider a *multiset* $\mathcal{M}$ of size n (e.g., $\mathcal{M} = \{1,2,3,3,4,5,6,7,8,9\}$ of size $n = 10$). The *multiplicity* of an element x of $\mathcal{M}$ is the number m_x of occurrences of x in $\mathcal{M}$ (e.g., $m_3 = 2$). We call the distribution of the multiplicities of the elements in $\mathcal{M}$ the *input structure* , and denote it by a set of pairs (x, m_x) (e.g., $\{(1,1), (2,1), (3,2), (4,1), (5,1), (6,1), (7,1), (8,1), (9,1)\}$ in $\mathcal{M}$). As early as 1976, Munro and Spira [25] described a variant of the algorithm `MergeSort` using counters, which optimally takes advantage of the input structure (i.e., the multiplicities of the distinct elements) when sorting a multiset $\mathcal{M}$ of n elements. Munro and Spira measure the "difficulty" of the instance in terms of the "input structure" by the entropy function $\mathcal{H}(m_1, \ldots, m_\sigma) = \sum_{i=1}^{\sigma} \frac{m_i}{n} \log \frac{n}{m_i}$. The time complexity of the algorithm is within $O(n(1 + \mathcal{H}(m_1, \ldots, m_\sigma))) \subseteq O(n(1+\log \sigma)) \subseteq O(n \log n)$, where σ is the number of distinct elements in $\mathcal{M}$ and $m_1, \ldots, m_\sigma$ are the multiplicities of the σ distinct elements in $\mathcal{M}$ (such that $\sum_{i=1}^{\sigma} m_i = n$), respectively.

Any array $\mathcal{A}$ representing a multiset lists its element in some order, which we call the *input order* and denote by a tuple (e.g., $\mathcal{A} = (2,3,1,3,7,8,9,4,5,6)$). Maximal sorted subblocks in $\mathcal{A}$ are a local form of input order and are called *runs* [21] (e.g., $\{(2,3), (1,3,7,8,9), (4,5,6)\}$ in $\mathcal{A}$). As early as 1973, Knuth [21] described a variant of the algorithm `MergeSort` using a prepossessing step taking linear time to detect *runs* in the array $\mathcal{A}$, which he named `Natural MergeSort`. Mannila [22] refined the analysis of the `Natural MergeSort` algorithm to yield a time complexity for sorting an array $\mathcal{A}$ of size n in time within $O(n(1+\log \rho)) \subseteq O(n \lg n)$, where ρ is the number of *runs* in the $\mathcal{A}$. Takaoka [28] described a new sorting algorithm that optimally takes

advantage of the distribution of the sizes of the runs in the array $\mathcal{A}$, which yields a time complexity within $O(n(1+\mathcal{H}(r_1,\ldots,r_\rho))) \subseteq O(n(1+\log\rho)) \subseteq O(n\log n)$, where ρ is the number of runs in $\mathcal{A}$ and $r_1,\ldots,r_\rho$ are the sizes of the ρ *runs* in $\mathcal{A}$ (such that $\sum_{i=1}^{\rho} r_i = n$), respectively. It is worth noting that in 1997, Takaoka [27] first described this algorithm in a technical report.

Given an element x of a multiset $\mathcal{M}$ and an integer $j \in [1..n]$ (e.g., $\mathcal{M} = \{1,2,3,3,4,5,6,7,8,9\}$, $x = 3$ and $j = 4$) , the *rank* `rank`(x) of x is the number of elements smaller than x in $\mathcal{M}$, (e.g., `rank`$(3) = 2$) and *selecting* the j-th element in $\mathcal{M}$ corresponds to computing the value `select`(j) of the j-th smallest element (counted with multiplicity) in $\mathcal{M}$ (e.g., `select`$(4) = 3$). As early as 1961, Hoare [17] showed how to support `rank` and `select` queries in average linear time, a result later improved to worst case linear time by Blum et al. [6]. Twenty years later, Dobkin and Munro [10] described a MultiSelection algorithm that supports several `select` queries in parallel and whose running time is optimal in the worst case over all multisets of size n and all sets of q queries hitting positions in the multisets separated by *gaps* (differences between consecutive `select` queries in sorted order) of sizes $g_0,\ldots,g_q$. Kaligosi et al. [19] later described a variant of this algorithm which number of comparisons performed is within a negligible additional term of the optimal. In the online context where the queries arrive one at a time, Karp et al. [20] further extended Dobkin and Munro's result [10]. Karp et al. called their solution a Deferred Data Structure and describe it as "lazy", as it partially sorts data, performing the minimum amount of work necessary in the worst case over all instances for a fixed n and q. Barbay et al. [2] refined this result by taking advantage of the gaps between the positions hit by the queries (i.e., the *query structure*). This suggests the following questions:

1. **Is there a sorting algorithm for multisets which takes the best advantage of both its *input order* and its *input structure* in a synergistic way, so that it performs as good as previously known solutions on all instances, and much better on instances where it can take advantage of both at the same time?**
2. **Is there a multiselection algorithm and a deferred data structure for answering `rank` and `select` queries which takes the best advantage not only of both of those notions of easiness in the input, but additionally also of notions of easiness in the queries, such as the *query structure* and the *query order*?**

We answer both questions affirmatively: In the context of Sorting, this improves upon both algorithms from Munro and Spira [25] and Takaoka [28]. In the context of MultiSelection and Deferred Data Structure for `rank` and `select` on Multisets, this improves upon Barbay et al.'s results [2] by adding 3 new measures of difficulty (input order, input structure and query order) to the single one previously considered (query structure). Even though the techniques used by our algorithms are known, the techniques used to refine the analysis of these algorithms to show that they improve the state of the art are complex. Additionally, we correct the analysis of the `Sorted Set Union` algorithm by Demaine et al. [9] (Section 2.3), and we define a simple yet new notion of "global" input order (Section 2.5), formed by the number of preexisting pivot positions in the input (e.g. $(3,2,1,6,5,4)$ has one pre-existing pivot position in the middle), not mentioned in previous surveys [11, 23] nor extensions [3].

We present our results incrementally, each building on the previous one, such that the most complete and complex result is in Section 4. In Section 2 we describe how to measure the interaction of the order (local and global) with the structure in the input, and two new synergistic Sorting algorithms based on distinct paradigms (i.e., merging vs splitting) which take advantage of both the input order and structure in order to sort the multiset in less time than traditional solutions based on one of those, at most. We refine the second of those results in Section 3 with the analysis of a MultiSelection algorithm which takes advantage of not only the order and structure in the input, but also of the *query structure*, in the offline setting. In Section 4 we analyze an online Deferred Data Structure taking advantage of the order and structure in the input on one hand, and of the order and structure in the queries on the other hand, in a synergistic way. As an additional result, we describe in Section 5 two compressed data structures to represent a multiset taking advantage of both the input order and structure, while supporting the operators `rank` and `select` on the multiset. We conclude with a discussion of our results in Section 6.

2 Sorting Algorithms

We review in Section 2.1 the algorithms `MergeSort with Counters` described by Munro and Spira [25] and `Minimal MergeSort` described by Takaoka [28], each takes advantage of distinct features in the input. In Section 2.2 we show that the algorithm `MergeSort with Counters` is incomparable with the algorithm `Minimal MergeSort`, in the sense that none performs always better than the other (by describing families of instances where `MergeSort with Counters` performs better than `Minimal MergeSort` on one hand, and families of instances where the second one performs better than the first one on the other hand), and some simple modifications and combinations of these algorithms, which still do not take full advantage of both the order (local and global) and structure in the input. In Sections 2.3 and 2.4, we describe two synergistic SORTING algorithms, which never perform worse than the algorithms presented in Section 2.1 and Section 2.2, and perform much better on some large classes of instances by taking advantage of both the order (local and global) and the structure in the input, in a synergistic way.

2.1 Previously Known Input-structure and Local Input-order Algorithms

The algorithm `MergeSort with Counters` described by Munro and Spira [25] is an adaptation of the traditional sorting algorithm `MergeSort` that optimally takes advantage of the structure in the input when sorting a multiset $\mathcal{M}$ of size n. The algorithm divides $\mathcal{M}$ into two equal size parts, sorts both parts recursively, and then merges the two sorted lists. When two elements of same value v are found, one is thrown away and a counter holding the number of occurrences of v is updated. Munro and Spira measure the "difficulty" of the instance in terms of the "input structure" by the entropy function $\mathcal{H}(m_1, \ldots, m_\sigma) = \sum_{i=1}^{\sigma} \frac{m_i}{n} \log \frac{n}{m_i}$. The time complexity of the algorithm is within $O(n(1 + \mathcal{H}(m_1, \ldots, m_\sigma))) \subseteq O(n(1+\log \sigma)) \subseteq O(n \log n)$, where σ is the number of distinct elements in $\mathcal{M}$ and $m_1, \ldots, m_\sigma$ are the multiplicities of the σ distinct elements in $\mathcal{M}$ (such that $\sum_{i=1}^{\sigma} m_i = n$), respectively.

The algorithm `Minimal MergeSort` described by Takaoka [28] optimally takes advantage of the local order in the input, as measured by the decomposition into runs when sorting an array $\mathcal{A}$ of size n. The main idea is to detect the runs first and then merge them pairwise, using a `MergeSort`-like step. The runs are detected in linear time by a scanning process identifying the positions i in $\mathcal{A}$ such that $\mathcal{A}[i] > \mathcal{A}[i+1]$. Merging the two shortest runs at each step further reduces the number of comparisons, making the running time of the merging process adaptive to the entropy of the sequence of the sizes of the runs. The merging process is then represented by a tree with the shape of a Huffman [18] tree, built from the distribution of the sizes of the *runs*. If the array $\mathcal{A}$ is formed by ρ runs and $r_1, \ldots, r_\rho$ are the sizes of the ρ runs (such that $\sum_{i=1}^{\rho} r_i = n$), then the algorithm sorts $\mathcal{A}$ in time within $O(n(1 + \mathcal{H}(r_1, \ldots, r_\rho))) \subseteq O(n(1+\log \rho)) \subseteq O(n \log n)$.

2.2 Comparison Between Sorting Algorithms

Example 1: $\langle 1, 2, 1, 2, \ldots, 1, 2 \rangle$

On this family of instances, the algorithm `Minimal MergeSort` detects $\frac{n}{2}$ runs that merges two at a time. The time complexity of `Minimal MergeSort` is within $\Theta(n \log n)$. On the other hand, the algorithm `MergeSort with Counters` recognizes the multiplicity of the elements with values 1 and 2. In every merging step, the resulting set is always $\{1, 2\}$. Therefore, the number of elements is reduced by half at each step, which yields a complexity linear in the size of the instance for sorting such instances.

Example 2: $\langle 1, 2, \ldots, n \rangle$

On this family of instances, the algorithm `Minimal MergeSort` detects in time linear in the size n of the input that the sequence is already sorted, and in turn it finishes. On the other hand, the running time of the algorithm `MergeSort with Counters` is within $\Theta(n \log n)$ because this algorithm would perform at least as many operations as `MergeSort`.

[4] The algorithm `Parallel Minimal Counter MergeSort` runs both algorithms in parallel, when one of these algorithms manages to sort the sequence, the algorithm `Parallel Minimal Counter MergeSort` returns the sorted sequence and then finishes. The time complexity of this algorithm in any instance is twice the minimum of the complexities of `Minimal MergeSort` and `MergeSort with Counters` for this instance. The algorithm `Parallel Minimal Counter MergeSort` needs to duplicate the input in order to run both algorithms in parallel.

Combining the ideas of identifying and merging runs from Takaoka [28] with the use of counters by Munro and Spira [25], we describe the `Small vs Small Sort` algorithm to sort a multiset. It identifies the runs using the same linear scanning as `Minimal MergeSort` and associates counters to the elements in the same way that `MergeSort with Counters` does. Once the runs are identified, this algorithm initializes a heap with them ordered by sizes. At each merging step the two shorter runs are selected for merging and both are removed from the heap. The pair is merged and the resulting run is inserted into the heap. The process is repeated until only one run is left and the sorted sequence is known. The `Small vs Small Sort` algorithm is adaptive to the sizes of the resulting runs in the merging process.

The time complexity of `Small vs Small Sort` for all instances is within a constant factor of the time complexity of `Parallel Minimal Counter MergeSort`; but the next example shows that there are families of instances where `Small vs Small Sort` performs better than a constant factor of the time complexity of `Parallel Minimal Counter MergeSort`.

Example 3: $\langle 1, 2, \ldots, \sigma, 1, 2, \ldots, \sigma \ldots, 1, 2, \ldots, \sigma \rangle$

In this family of instances, there are ρ runs, each of size σ. The complexity of `MergeSort with Counters` for this instance is within $\Theta(\rho\sigma \log \sigma)$, while the complexity of `Minimal MergeSort` for this instance is within $\Theta(\rho\sigma \log \rho)$. On the other hand, the complexity `Small vs Small Sort` for this instance is better, within $\Theta(\rho\sigma)$: at each level of the binary tree representing the merging order, the sum of the sizes of the runs is halved.

Even though `Small vs Small Sort` is adaptive to the sizes of the resulting runs, it does not take advantage of the fact that there may exist a pair of runs that can be merged very quickly, but it needs linear time in the sum of the sizes to merge one of them when paired with another run of the same size. We show this disadvantage in the Example 4 at the end of the next section.

In the following sections we describe two sorting algorithms that take the best advantage of both the order (local and global) and structure in the input all at once when sorting a multiset. The first one is a straightforward application of previous results, while the second one prepares the ground for the MultiSelection algorithm (Section 3) and the Deferred Data Structures (Section 4), which take advantage of the order (local and global) and structure in the data and of the order and structure in the queries.

2.3 "Kind-of-new" Sorting Algorithm `DLM Sort`

In 2000, Demaine et al. [9] described the algorithm `DLM Union`, an instance optimal algorithm that computes the union of ρ sorted sets. It inserts the smallest element of each set in a heap. At each step, it deletes from the heap all the elements whose values are equal to the minimum value of the heap. If more than one element is deleted, it knows the multiplicity of this value in the union of the sets. It then adds to the heap the elements following the elements of minimum value of each set that contained the minimum value. If there is only one minimum element, it extracts from the heap the second minimum and executes a doubling search [5] in the set where the minimum belongs for the value of the second minimum. Once it finds the insertion rank r of the second minimum (i.e., number of elements smaller than the second minimum in the set that contains the minimum), it also knows that the multiplicity of all elements whose positions are before r in the set that contain the minimum are 1 in the union of the ρ sets. The process is repeated until all elements are discarded.

The time complexity of the `DLM Union` algorithm is measured in terms of the number and sizes of blocks of consecutive elements in the sets that are also consecutive in the sorted union (see Figure 1 for a graphical representation of such a decomposition on a particular instance of the Sorted Set Union problem). The

sizes of these blocks are referred to as *gaps* in the analysis of the algorithm. These blocks induce a partition π of the output into intervals such that any singleton corresponds to a value that has multiplicity greater than 1 in the input, and each other interval corresponds to a block as defined above. Each member i of π has a value m_i associated with it: if the member i of π is a block, then m_i is 1, otherwise, if the member i of π is a singleton corresponding to a value of multiplicity q, then m_i is q. Let χ be the size of π. If the instance is formed by δ blocks of sizes $g_1, \ldots, g_\delta$ such that these blocks induce a partition π whose members have values $m_1, \ldots, m_\chi$, we express the time complexity of `DLM Union` as within $\Theta(\sum_{i=1}^{\delta} \log g_i + \sum_{i=1}^{\chi} \log \binom{\rho}{m_i})$. This time complexity is within a constant factor of the complexity of any other algorithm computing the union of these sorted sets (i.e., the algorithm is instance optimal).

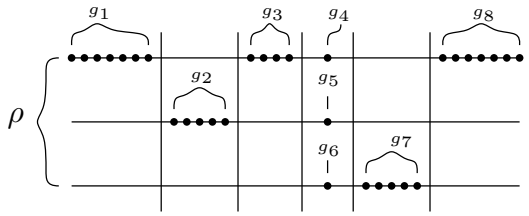

Fig. 1: An instance of the Sorted Set Union problem with $\rho = 3$ sorted sets. In each set, the entry $\mathcal{A}[i]$ is represented by a point of x-coordinate $\mathcal{A}[i]$. The blocks that form the sets are noted. The blocks g_4, g_5 and g_6 are of size 1 because they correspond to elements of equal value and they induce the 4-th member of the partition π with value m_4 equals 3. The vertical bars separate the members of π.

We adapt the `DLM Union` algorithm for sorting a multiset. The algorithm `DLM Sort` detects the runs first through a linear scan and then applies the algorithm `DLM Union`. After that, transforming the output of the union algorithm to yield the sorted multiset takes only linear time. The following corollary follows from our refined analysis above:

Corollary 1. *Given a multiset $\mathcal{M}$ of size n formed by ρ runs and δ blocks of sizes $g_1, \ldots, g_\delta$ such that these blocks induce a partition π of size χ of the output whose members have values $m_1, \ldots, m_\chi$, the algorithm* `DLM Sort` *performs within $O(n + \sum_{i=1}^{\delta} \log g_i + \sum_{i=1}^{\chi} \log \binom{\rho}{m_i})$ data comparisons. This number of comparisons is optimal in the worst case over multisets of size n formed by ρ runs and δ blocks of sizes $g_1, \ldots, g_\delta$ such that these blocks induce a partition π of size χ of the output whose members have values $m_1, \ldots, m_\chi$.*

There are families of instances where the algorithm `DLM Sort` behaves significantly better than the algorithm `Small vs Small Sort`. Consider for instance the following example:

Example 4: $\left\langle \frac{\rho-1}{\rho}n + 1, \ldots, n, \frac{\rho-2}{\rho}n + 1, \ldots, \frac{\rho-1}{\rho}n, \ldots, 1, \ldots, \frac{1}{\rho}n \right\rangle$

In this family of instances, there are ρ runs of size $\frac{n}{\rho}$ each. The runs are pairwise disjoint and the elements of each run are consecutive in the sorted set. The time complexity of the algorithm `Small vs Small Sort` in this instances is within $\Theta(n \log \rho)$, while the time complexity of the algorithm `DLM Sort` is within $\Theta\left(n + \rho \log \rho + \rho \log \frac{n}{\rho}\right) = \Theta(n + \rho \log n)$ (which is better than $\Theta(n \log \rho)$ for $\rho \in o(n)$).

While the algorithm `DLM Sort` answers the Question 1 from Section 1, it does not yield a MultiSelection algorithm nor a Deferred Data Structure answering Question 2. In the following section we describe another sorting algorithm that also optimally takes advantage of the local order and structure in the input, but which is based on a distinct paradigm, more suitable to such extensions.

2.4 New Sorting Algorithm `Quick Synergy Sort`

Given a multiset $\mathcal{M}$, the algorithm `Quick Synergy Sort` identifies the *runs* in linear time through a scanning process. As indicated by its name, the algorithm is directly inspired by the `QuickSort` [16] algorithm. It computes a pivot μ, which is the median of the set formed by the middle elements of each run, and partitions each *run* by the value of μ. This partitioning process takes advantage of the fact that the elements in each *run* are already sorted. It then recurses on the elements smaller than μ and on the elements greater than μ. (See Algorithm 1 for a more formal description).

Definition 1 (Median of the middles). *Given a multiset $\mathcal{M}$ formed by runs, the "*median of the middles*" is the median element of the set formed by the middle elements of each run.*

Algorithm 1 `Quick Synergy Sort`

Input: A multiset $\mathcal{M}$ of size n
Output: A sorted sequence of $\mathcal{M}$

1: Compute the ρ runs of respective sizes $(r_i)_{i\in[1..\rho]}$ in $\mathcal{M}$ such that $\sum_{i=1}^{\rho} r_i = n$;
2: Compute the median μ of the middles of the runs, note $j \in [1..\rho]$ the run containing μ;
3: Perform doubling searches for the value μ in all runs except the j-th, starting at both ends of the runs in parallel;
4: Find the maximum $\max_\ell$ (minimum $\min_r$) among the elements smaller (resp., greater) than μ in all runs except the j-th;
5: Perform doubling searches for the values $\max_\ell$ and $\min_r$ in the j-th run, starting at the position of μ;
6: Recurse on the elements smaller than o equal to $\max_\ell$ and on the elements greater than or equal to $\min_r$.

The number of data comparisons performed by the algorithm `Quick Synergy Sort` is asymptotically the same as the number of data comparisons performed by the algorithm `DLM Sort` described in the previous section. We divide the proof into two lemmas. We first bound the overall number of data comparisons performed by all the doubling searches of the algorithm `Quick Synergy Sort`.

Lemma 1. *Let $g_1, \ldots, g_k$ be the sizes of the k blocks that form the r-th run. The overall number of data comparisons performed by the doubling searches of the algorithm `Quick Synergy Sort` to find the values of the medians of the middles in the r-th run is within $O(\sum_{i=1}^{k} \log g_i)$.*

Proof. Every time the algorithm finds the insertion rank of one of the medians of the middles in the r-th run, it partitions the run by a position separating two blocks. The doubling search steps can be represented as a tree (see Figure 2 for a tree representation of a particular instance). Each node of the tree corresponds to a step. Each internal node has two children, which correspond to the two subproblems into which the step partitions the run. The cost of the step is less than four times the logarithm of the size of the child subproblem with smaller size, because of the two doubling searches in parallel. The leaves of the tree correspond to the blocks themselves.

We prove that at each step, the total cost is bounded by eight times the sum of the logarithms of the sizes of the leaf subproblems. This is done by induction over the number of steps. If the number of steps is zero then there is no cost. For the inductive step, if the number of steps increases by one, then a new doubling search step is done and a leaf subproblem is partitioned into two new subproblems. At this step, a leaf of the tree is transformed into an internal node and two new leaves are created. Let a and b such that $a \leq b$ be the sizes of the new leaves created. The cost of this step is less than $4 \log a$. The cost of all the steps then increases by $4 \lg a$, and hence the sum of the logarithms of the sizes of the leaves increases by $8(\lg a + \lg b) - 8 \lg(a + b)$. But if $a \geq 4$ and $b \geq a$, then $2 \lg(a + b) \leq \lg a + 2 \lg b$. The result follows. □

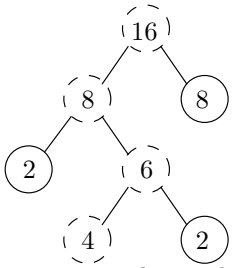

Fig. 2: The tree that represents the doubling search steps for a run composed of four blocks of respective sizes $2, 4, 2, 8$. The size of the subproblem is noted in each node. At each subproblem, the cost of the step is the logarithm of the size of the subproblem of the solid child.

The step that computes the median μ of the middles of ρ runs and the step that finds the maximum $\max_\ell$ (minimum $\min_r$) among the elements smaller (resp., greater) than μ of ρ runs performs linear in ρ data

comparisons. As shown in the following lemma, the overall number of data comparisons performed during these steps is within $O(\sum_{i=1}^{\chi} \log \binom{\rho}{m_i})$, where $m_1, \ldots, m_\chi$ are the values of the member of the partition π (see Section 2.3 for the definition of π) and ρ is the number of runs in $\mathcal{M}$.

Consider the instance depicted in Figure 3 for an example illustrating where the term $\log \binom{\rho}{m_v}$ comes from. In this instance, there is a value v that has multiplicity $m_v > 1$ in $\mathcal{M}$ and the rest of the values have multiplicity 1. The elements with value v are present at the end of the last m_v runs and the rest of the runs are formed by only one block. The elements of the i-th run are greater than the elements of the $(i+1)$-th run. During the computation of the medians of the middles, the number of data comparisons that involve elements of value v is within $O(\log \binom{\rho}{m_v})$. The algorithm computes the median μ of the middles and partitions the runs by the value of μ. In the recursive call that involves elements of value v, the number of runs is reduced by half. This is repeated until one occurrence of μ belongs to one of the last m_v runs. The number of data comparisons that involve elements of value v up to this step is within $O(m_v \log \frac{\rho}{m_v}) = O(\log \binom{\rho}{m_v})$, where $\log \frac{\rho}{m_v}$ corresponds to the number of steps where μ does not belong to the last m_v runs. The next recursive call will necessarily choose one element of value v as the median of the middles.

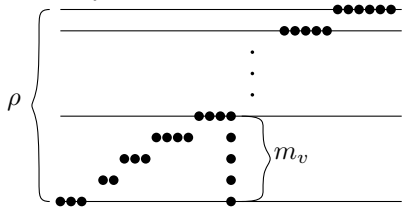

Fig. 3: A multiset $\mathcal{M}$ formed by ρ runs. Each entry $\mathcal{M}[i]$ is represented by a point of x-coordinate $\mathcal{M}[i]$. There is an element of multiplicity m_v present in the last m_v runs and the rest of the runs are formed by only one block.

Lemma 2. *Let $\mathcal{M}$ be a multiset formed by ρ runs and δ blocks such that these blocks induce a partition π of the output of size χ whose members have values $m_1, \ldots, m_\chi$. Consider the steps that compute the medians of the middles and the steps that find the elements* $\max_\ell$ *and* $\min_r$ *in the algorithm* ***Quick Synergy Sort****, the overall number of data comparisons performed during these steps is within $O(\sum_{i=1}^{\chi} \log \binom{\rho}{m_i})$.*

Proof. We prove this lemma by induction over the size of π and ρ. The number of data comparisons performed by one of these steps is linear in the number of runs in the sub-instance (i.e., ignoring all the empty sets of this sub-instance). Let $\mathcal{T}(\pi, \rho)$ be the overall number of data comparisons performed during the steps that compute the medians of the middles and during the steps that find the elements $\max_\ell$ and $\min_r$ in the algorithm `Quick Synergy Sort`. We prove that $\mathcal{T}(\pi, \rho) \leq \sum_{i=1}^{\chi} m_i \log \frac{\rho}{m_i} - \rho$. Let μ be the first median of the middles computed by the algorithm. Let ℓ and r be the number of runs that are completely to the left and to the right of μ, respectively. Let b be the number of runs that are split in the doubling searches for the value μ in all runs. Let π_ℓ and π_r be the partitions induced by the blocks yielded to the left and to the right of μ, respectively. Then, $\mathcal{T}(\pi, \rho) = \mathcal{T}(\pi_\ell, \ell + b) + \mathcal{T}(\pi_r, r + b) + \rho$ because of the two recursive calls and the step that computes μ. By Induction Hypothesis, $\mathcal{T}(\pi_\ell, \ell + b) \leq \sum_{i=1}^{\chi_\ell} m_i \log \frac{\ell+b}{m_i} - \ell - b$ and $\mathcal{T}(\delta_r, r + b) \leq \sum_{i=1}^{\chi_r} m_i \log \frac{r+b}{m_i} - r - b$. Hence, we need to prove that $\ell + r \leq \sum_{i=1}^{\chi_\ell} m_i \log \left(1 + \frac{r}{\ell+b}\right) + \sum_{i=1}^{\chi_r} m_i \left(1 + \frac{\ell}{r+b}\right)$, but this is a consequence of $\sum_{i=1}^{\chi_\ell} m_i \geq \ell + b, \sum_{i=1}^{\chi_r} m_i \geq r + b$ (the number of blocks is greater than or equal to the number of runs); $\ell \leq r + b, r \leq \ell + b$ (at least $\frac{\rho}{2}$ runs are left to the left and to the right of μ); and $\log \left(1 + \frac{y}{x}\right)^x \geq y$ for $y \leq x$. □

Consider the step that performs doubling searches for the values $\max_\ell$ and $\min_r$ in the run that contains the median μ of the middles, this step results in the finding of the block g that contains μ in at most $4 \log |g|$ data comparisons, where $|g|$ is the size of g. Combining Lemma 1 and Lemma 2 yields an upper bound on the number of data comparisons performed by the algorithm `Quick Synergy Sort`:

Theorem 1. *Let $\mathcal{M}$ be a multiset of size n formed by ρ runs and δ blocks of sizes $g_1, \ldots, g_\delta$ such that these blocks induce a partition π of the output of size χ whose members have values $m_1, \ldots, m_\chi$. The algorithm* ***Quick Synergy Sort*** *performs within $O(n + \sum_{i=1}^{\delta} \log g_i + \sum_{i=1}^{\chi} \log \binom{\rho}{m_i})$ data comparisons on $\mathcal{M}$. This number of comparisons is optimal in the worst case over multisets of size n formed by ρ runs and δ blocks of sizes $g_1, \ldots, g_\delta$ such that these blocks induce a partition π of size χ of the output whose members have values $m_1, \ldots, m_\chi$.*

 We extend these results to take advantage of the global order of the multiset in a way that can be combined with the notion of runs (local order).

2.5 Taking Advantage of Global Order

Given a multiset $\mathcal{M}$, a *pivot position* is a position p in $\mathcal{M}$ such that all elements in previous position are smaller than or equal to all elements at p or in the following positions. Formally:

Definition 2 (Pivot positions). *Given a multiset $\mathcal{M} = (x_1, \ldots, x_n)$ of size n, the pivot positions are the positions p such that $x_a \leq x_b$ for all a, b such that $a \in [1..p-1]$ and $b \in [p..n]$.*

Existing pivot positions in the input order of $\mathcal{M}$ divide the input into subsequences of consecutive elements such that the range of positions of the elements at each subsequence coincide with the range of positions of the same elements in the sorted sequence of $\mathcal{M}$: the more there are of such positions, the more "global" order there is in the input. Detecting such positions takes only a linear number of comparisons.

Lemma 3. *Given a multiset $\mathcal{M}$ of size n with ϕ pivot positions $p_1, \ldots, p_\phi$, the ϕ pivot positions can be detected within a linear number of comparisons.*

Proof. The `bubble-up` step of the algorithm `BubbleSort` [21] sequentially compares the elements in positions $i-1$ and i of $\mathcal{M}$, for i from 2 to n. If $\mathcal{M}[i-1] > \mathcal{M}[i]$, then the elements interchange their values. As consequence of this step the elements with large values tend to move to the right. In an execution of a `bubble-up` step in $\mathcal{M}$, the elements that do not interchange their values are those elements whose values are greater than or equal to all the elements on their left. The `bubble-down` step is similar to the `bubble-up` step, but it scans the sequence from right to left, interchanging the elements in positions $i-1$ and i if $\mathcal{M}[i-1] > \mathcal{M}[i]$. In an execution of a `bubble-down` step in $\mathcal{M}$, the elements that do not interchange their values are those elements whose values are smaller than or equal to all the elements on their right. Hence, the positions of the elements that do not interchange their values during the executions of both `bubble-up` and `bubble-down` steps are the pivot positions in $\mathcal{M}$. □

When there are ϕ such positions, they simply divide the input of size n into $\phi+1$ sub-instances of sizes $n_0, \ldots, n_\phi$ (such that $\sum_{i=0}^{\phi} n_i = n$). Each sub-instance I_i for $i \in [0..\phi]$ then has its own number of runs r_i and alphabet size σ_i, on which the synergistic solutions described in this work can be applied, from mere Sorting (Section 2) to supporting MultiSelection (Section 3) and the more sophisticated Deferred Data Structures (Section 4).

Corollary 2. *Let $\mathcal{M}$ be a multiset of size n with ϕ pivot positions. The ϕ pivot positions divide $\mathcal{M}$ into $\phi+1$ sub-instances of sizes $n_0, \ldots, n_\phi$ (such that $\sum_{i=0}^{\phi} n_i = n$). Each sub-instance I_i of size n_i is formed by ρ_i runs and δ_i blocks of sizes $g_{i1}, \ldots, g_{i\delta_i}$ such that these blocks induce a partition π_i of the output of size χ_i whose members have values $m_{i1}, \ldots, m_{i\chi_i}$ for $i \in [0..\phi]$. There exists an algorithm that performs within $O(n + \sum_{i=0}^{\phi} \left\{ \sum_{j=1}^{\delta_i} \log g_{ij} + \sum_{j=1}^{\chi_i} \log \binom{\rho}{m_{ij}} \right\})$ data comparisons for sorting $\mathcal{M}$. This number of comparisons is optimal in the worst case over multisets of size n with ϕ pivot positions which divide the multiset into $\phi+1$ sub-instances of sizes $n_0, \ldots, n_\phi$ (such that $\sum_{i=0}^{\phi} n_i = n$) and each sub-instance I_i of size n_i is formed by ρ_i runs and δ_i blocks of sizes $g_{i1}, \ldots, g_{i\delta_i}$ such that these blocks induce a partition π_i of the output of size χ_i whose members have values $m_{i1}, \ldots, m_{i\chi_i}$ for $i \in [0..\phi]$.*

Next, we generalize the algorithm `Quick Synergy Sort` to an offline multiselection algorithm that partially sorts a multiset according to the set of `select` queries given as input. This serves as a pedagogical introduction to the online Deferred Data Structures for answering `rank` and `select` queries presented in Section 4.

For simplicity, in Section 3 and Section 4 we first describe results ignoring existing pivot positions, and then present the complete result as corollary.

3 MultiSelection Algorithm

Given a linearly ordered multiset $\mathcal{M}$ and a sequence of ranks $r_1, \ldots, r_q$, a multiselection algorithm must answer the queries `select`(r_1), ..., `select`(r_q) in $\mathcal{M}$, hence partially sorting $\mathcal{M}$. We describe a MultiSelection algorithm based on the sorting algorithm `Quick Synergy Sort` introduced in Section 2.4. This algorithm is an intermediate result leading to the two Deferred Data Structures described in Section 4.

Given a multiset $\mathcal{M}$ and a set of q `select` queries, the algorithm `Quick Synergy MultiSelection` follows the same first steps as the algorithm `Quick Synergy Sort`. But once it has computed the ranks of all elements in the block that contains the pivot μ, it determines which `select` queries correspond to elements smaller than or equal to $\max_\ell$ and which ones correspond to elements greater than or equal to $\min_r$ (see Algorithm 1 for the definitions of $\max_\ell$ and $\min_r$). It then recurses on both sides. See Algorithm 2 for a formal description of the algorithm `Quick Synergy MultiSelection`.

Algorithm 2 `Quick Synergy MultiSelection`

Input: A multiset $\mathcal{M}$ and a set Q of q offline `select` queries
Output: The q selected elements

1: Compute the ρ runs of respective sizes $(r_i)_{i\in[1..\rho]}$ in $\mathcal{M}$ such that $\sum_{i=1}^{\rho} r_i = n$;
2: Compute the median μ of the middles of the ρ runs, note $j \in [1..\rho]$ the run containing μ;
3: Perform doubling searches for the value μ in all runs except the j-th, starting at both ends of the runs in parallel;
4: Find the maximum $\max_\ell$ (minimum $\min_r$) among the elements smaller (resp., greater) than μ in all runs except the j-th;
5: Perform doubling searches for the values $\max_\ell$ and $\min_r$ in the j-th run, starting at the position of μ;
6: Compute the set of queries Q_ℓ that go to the left of $\max_\ell$ and the set of queries Q_r that go to the right of $\min_r$;
7: Recurse on the elements smaller than or equal to $\max_\ell$ and on the elements greater than or equal to $\min_r$ with the set of queries Q_ℓ and Q_r, respectively.

We extend the notion of blocks to the context of partial sorting. The idea is to consider consecutive blocks, which have not been identified by the `Quick Synergy MultiSelection` algorithm, as a single block. We next introduce the definitions of *pivot blocks* and *selection blocks*.

Definition 3 (Pivot Blocks). *Given a multiset $\mathcal{M}$ formed by ρ runs and δ blocks. The "*pivot blocks*" are the blocks of $\mathcal{M}$ that contain the pivots and the elements of value equals to the pivots during the steps of the algorithm* `Quick Synergy MultiSelection`*.*

In each run, between the pivot blocks and the insertion ranks of the pivots, there are consecutive blocks that the algorithm `Quick Synergy MultiSelection` has not identified as separated blocks, because no doubling searches occurred inside them.

Definition 4 (Selection Blocks). *Given the i-th run, formed of various blocks, and q* `select` *queries, the algorithm* `Quick Synergy MultiSelection` *computes ξ pivots in the process of answering the q queries. During the doubling searches, the algorithm* `Quick Synergy MultiSelection` *finds the insertion ranks of the ξ pivots inside the i-th run. These positions determine a partition of size $\xi+1$ of the i-th run where each element of the partition is formed by consecutive blocks or is empty. We call the elements of this partition "*selection blocks*". The set of all selection blocks include the set of all pivot blocks.*

Using these definitions, we generalize the results proven in Section 2.4 to the more general problem of MultiSelection.

Theorem 2. *Given a multiset $\mathcal{M}$ of size n formed by ρ runs and δ blocks; and q offline* `select` *queries over $\mathcal{M}$ corresponding to elements of* `ranks` *$r_1, \ldots, r_q$. The algorithm* `Quick Synergy MultiSelection` *computes ξ pivots in the process of answering the q queries. Let $s_1, \ldots, s_\beta$ be the sizes of the β selection*

blocks determined by these ξ pivots in all runs. Let $m_1, \ldots, m_\lambda$ be the numbers of pivot blocks among this selection blocks corresponding to the values of the λ pivots with multiplicity greater than 1, respectively. Let $\rho_0, \ldots, \rho_\xi$ be the sequence where ρ_i is the number of runs that have elements with values between the pivots i and $i+1$ sorted by `ranks`*, for $i \in [1..\xi]$. The algorithm* `Quick Synergy MultiSelection` *answers the q* `select` *queries performing within $O\left(n + \sum_{i=1}^{\beta} \log s_i + \beta \log \rho - \sum_{i=1}^{\lambda} m_i \log m_i - \sum_{i=0}^{\xi} \rho_i \log \rho_i\right) \subseteq O\left(n \log n - \sum_{i=0}^{q} \Delta_i \log \Delta_i\right)$ data comparisons, where $\Delta_i = r_{i+1} - r_i$, $r_0 = 0$ and $r_{q+1} = n$.*

Proof. The doubling searches that find the insertion ranks of the pivots during the overall execution of the algorithm perform within $O(\sum_{i=1}^{\beta} \log s_i)$ data comparisons . At each run, a constant factor of the sum of the logarithm of the sizes of the selection blocks bounds the number of data comparisons performed by these doubling searches (see the proof of Lemma 1 analyzing the algorithm `Quick Synergy Sort` for details).

The pivots computed by the algorithm `Quick Synergy MultiSelection` for answering the queries are a subset of the pivots computed by the algorithm `Quick Synergy Sort` for sorting the whole multiset. Suppose that the selection blocks determined by every two consecutive pivots form a multiset $\mathcal{M}_j$ such that for every pair of selection blocks in $\mathcal{M}_j$ the elements of one are smaller than the elements of the other one. Consider the steps that compute the medians of the middles in the algorithm `Quick Synergy Sort`, the number of data comparisons performed by these steps would be within $O\left(n + \sum_{i=1}^{\beta} \log s_i + \beta \log \rho - \sum_{i=1}^{\lambda} m_i \log m_i\right)$ in this supposed instance (see the proof of Lemmas 2 analyzing the algorithm `Quick Synergy Sort` for details). The number of comparisons needed to sort the multisets $\mathcal{M}_j$ is within $\Theta(\sum_{i=0}^{\xi} \rho_i \log \rho_i)$. The result follows. □

The process of detecting the ϕ pre-existing pivot positions seen in Section 2.5 can be applied as the first step of the multiselection algorithm. The ϕ pivot positions divide the input of size n into $\phi+1$ sub-instances of sizes $n_0, \ldots, n_\phi$. For each sub-instance I_i for $i \in [0..\phi]$, the multiselection algorithm determines which `select` queries correspond to I_i and applies then the steps of the Algorithm 2 inside I_i in order to answer these queries.

Corollary 3. *Let $\mathcal{M}$ be a multiset of size n with ϕ pivot positions. The ϕ pivot positions divide $\mathcal{M}$ into $\phi+1$ sub-instances of sizes $n_0, \ldots, n_\phi$ (such that $\sum_{i=0}^{\phi} n_i = n$). Let q be the number of offline* `select` *queries over $\mathcal{M}$, such that q_i queries correspond to the sub-instance I_i, for $i \in [0..\phi]$. In each sub-instance I_i of size n_i formed by ρ_i runs, the algorithm* `Quick Synergy MultiSelection` *selects ξ_i pivots when it answers the q_i queries. These ξ_i pivots determine β_i selection blocks of sizes $s_{i1}, \ldots, s_{i\beta_i}$ inside I_i. Let $m_{i1}, \ldots, m_{i\lambda_i}$ be the numbers of pivot blocks among this selection blocks corresponding to the values of the λ_i pivots with multiplicity greater than 1, respectively. Let $\rho_{i0}, \ldots, \rho_{i\xi_i}$ be the sequence where ρ_{ij} is the number of runs that have elements with values between the pivots ij and $i(j+1)$ sorted by* `ranks`*, for $j \in [1..\xi_i]$. There is an algorithm that answers the q offline* `select` *queries performing within $O(n + \sum_{i=0}^{\phi} \left\{\sum_{j=1}^{\beta_i} \log s_{ij} + \beta_i \log \rho_i - \sum_{j=1}^{\lambda_i} m_{ij} \log m_{ij} - \sum_{j=0}^{\xi_i} \rho_{ij} \log \rho_{ij}\right\})$ data comparisons.*

In the result above, the queries are given all at the same time (i.e., offline). In the context where they arrive one at the time (i.e., online), we define two Deferred Data Structures for answering online `rank` and `select` queries, both inspired by the algorithm `Quick Synergy MultiSelection`.

4 `Rank` and `Select` Deferred Data Structures

We describe two Deferred Data Structures that answer a set of `rank` and `select` queries arriving one at the time over a multiset $\mathcal{M}$, progressively sorting $\mathcal{M}$. Both data structures take advantage of the order (local and global) and structure in the input, and of the structure in the queries. The first data structure is in the RAM model of computation, at the cost of not taking advantage of the order in which the queries are given. The second data structure is in the comparison model (a more constrained model) but does take advantage of the query order.

4.1 Taking Advantage of Order and Structure in the Input, but only of Structure in the Queries

Given a multiset $\mathcal{M}$ of size n, the RAM Deferred Data Structure is composed of a bitvector $\mathcal{A}$ of size n, in which we mark the elements in $\mathcal{M}$ that have been computed as pivots by the algorithm when it answers the online queries; a dynamic predecessor and successor structure $\mathcal{B}$ over the bitvector $\mathcal{A}$, which allows us to find the two successive pivots between which the query fits; and for each pivot p found, the data structure stores pointers to the insertion ranks of p in each run, to the beginning and end of the block g to which p belongs, and to the position of p inside g. The dynamic predecessor and successor structure $\mathcal{B}$ requires the RAM model of computation in order to answer *predecessor and successor queries* in time within $o(\log n)$ [4].

Theorem 3. *Consider a multiset $\mathcal{M}$ of size n formed by ρ runs and δ blocks. The* RAM Deferred Data Structure *computes ξ pivots in the process of answering q online* `rank` *and* `select` *queries over $\mathcal{M}$. Let $s_1, \ldots, s_\beta$ be the sizes of the β selection blocks determined by these ξ pivots in all runs. Let $m_1, \ldots, m_\lambda$ be the numbers of pivot blocks among this selection blocks corresponding to the values of the λ pivots with multiplicity greater than 1, respectively. Let $\rho_0, \ldots, \rho_\xi$ be the sequence where ρ_i is the number of runs that have elements with values between the pivots i and $i+1$ sorted by* `ranks`*, for $i \in [1..\xi]$. Let u and $g_1, \ldots, g_u$ be the number of* `rank` *queries and the sizes of the identified and searched blocks in the process of answering the u* `rank` *queries, respectively. The* RAM Deferred Data Structure *answers these q online* `rank` *and* `select` *queries in time within $O(n + \sum_{i=1}^{\beta} \log s_i + \beta \log \rho - \sum_{i=1}^{\lambda} m_i \log m_i - \sum_{i=0}^{\xi} \rho_i \log \rho_i + \xi \log\log n + u \log n \log\log n + \sum_{i=1}^{u} \log g_i)$.*

Proof. The algorithm answers a new `select`(i) query by accessing in $\mathcal{A}$ the query position i. If $\mathcal{A}[i]$ is 1, then the element e has been computed as pivot, and hence the algorithm answers the query in constant time by following the position of e inside the block at which e belongs. If $\mathcal{A}[i]$ is 0, then the algorithm finds the nearest pivots to its left and right using the predecessor and successor structure, $\mathcal{B}$. If the position i is inside a block to which one of the two nearest pivots belong, then the algorithm answers the query and in turn finishes. If not, it then applies the same steps as the algorithm `Quick Synergy MultiSelection` in order to answer the query; it updates the bitvector $\mathcal{A}$ and the dynamic predecessor and successor structure $\mathcal{B}$ whenever a new pivot is computed; and for each pivot p computed, the structure stores the pointers to the insertion ranks of p in each run, to the beginning and end of the block g to which p belongs, and to the position of p inside g.

The algorithm answers a new `rank`(x) query by finding the *selection block* s_j in the j-th run such that x is between the smallest and the greatest value of s_j for all $j \in [1..\rho]$. For that, the algorithm performs a sort of parallel binary searches for the value x at each run taking advantage of the pivots that have been computed by the algorithm. The algorithm accesses the position $\frac{n}{2}$ in $\mathcal{A}$. If $\mathcal{A}[\frac{n}{2}]$ is 1, then the element e of `rank` $\frac{n}{2}$ has been computed as pivot. Following the pointer to the block g to which e belongs, the algorithm decides if x is to the right, to the left or inside g by performing a constant number of data comparisons. In the last case, a binary search for the value x inside g yields the answer of the query. If $\mathcal{A}[\frac{n}{2}]$ is 0, then the algorithm finds the nearest pivots to the left and right of the position $\frac{n}{2}$ using the predecessor and successor structure, $\mathcal{B}$. Following the pointers to the blocks that contain these pivots the algorithm decides if x is inside one of these blocks, to the right of the rightmost block, to the left of the leftmost block, or between these two blocks. In the last case, the algorithm applies the same steps as the algorithm `Quick Synergy MultiSelection` in order to compute the median μ of the middles and partitions the selection blocks by μ. The algorithm then decides to which side x belongs. These steps identify several new pivots, and in consequence several new blocks in the structure. □

The RAM Deferred Data Structure includes the pivot positions (seen in Section 2.5) as a natural extension of the algorithm. The ϕ pivot positions are marked in the bitvector $\mathcal{A}$. For each pivot position p, the structure stores pointers to the end of the runs detected on the left of p; to the beginning of the runs detected on the right of p; and to the position of p in the multiset.

Corollary 4. *Let $\mathcal{M}$ be a multiset of size n with ϕ pivot positions. The ϕ pivot positions divide $\mathcal{M}$ into $\phi+1$ sub-instances of sizes $n_0, \ldots, n_\phi$ (such that $\sum_{i=0}^{\phi} n_i = n$). Let q be the number of online* `rank` *and* `select`

queries over $\mathcal{M}$, such that q_i queries correspond to the sub-instance I_i, for $i \in [0..\phi]$. In each sub-instance I_i of size n_i formed by ρ_i runs, the RAM Deferred Data Structure *selects ξ_i pivots in the process of answering the q_i online* **rank** *and* **select** *queries over I_i. Let $s_{i1}, s_{i2}, \ldots, s_{i\beta_i}$ be the sizes of the β_i selection blocks determined by the ξ_i pivots in all runs of I_i. Let $m_{1i}, \ldots, m_{i\lambda_i}$ be the numbers of pivot blocks among this selection blocks corresponding to the values of the λ_i pivots with multiplicity greater than 1, respectively. Let $\rho_{i0}, \ldots, \rho_{i\xi_i}$ be the sequence where ρ_{ij} is the number of runs that have elements with values between the pivots ij and $i(j+1)$ sorted by* **ranks**, *for $j \in [1..\xi_i]$. Let u_i and $g_{i1}, \ldots, g_{iu_i}$ be the number of* **rank** *queries and the sizes of the identified and searched blocks in the process of answering the u_i* **rank** *queries over I_i, respectively. There exists an algorithm that answers these q online* **rank** *and* **select** *queries in time within $O(n + \sum_{i=0}^{\phi} \left\{ \beta_i \log \rho_i - \sum_{j=1}^{\lambda_i} m_{ij} \log m_{ij} - \sum_{j=0}^{\xi_i} \rho_{ij} \log \rho_{ij} + \xi_i \log\log n_i + u \log n_i \log\log n_i + \sum_{j=1}^{u_i} \log g_{ij} \right\})$.*

The RAM Deferred Data Structure takes advantage of the structure in the queries and of the structure and order (local and global) in the input. Changing the order in the `rank` and `select` queries does not affect the time complexity of the RAM Deferred Data Structure. Once the structure identifies the nearest pivots to the left and right of the query positions, the steps of the algorithms are the same as in the offline case (Section 3). We next describe a deferred data structure taking advantage of the structure and order in the queries and of the structure and order (local and global) in the input data.

4.2 Taking Advantage of the Order and Structure in both the Input and the Queries

To take advantage of the order in the queries, we introduce a data structure that finds the nearest pivots to the left and to the right of a position $p \in [1..n]$, while taking advantage of the distance between the position of the last computed pivot and p. This distance is measured in the number of computed pivots between the two positions. For that we use a *finger search tree* [15] which is a search tree maintaining *fingers* (i.e., pointers) to elements in the search tree. Finger search trees support efficient updates and searches in the vicinity of the fingers. Brodal [7] described an implementation of finger search trees that searches for an element x, starting the search at the element given by the finger f in time within $O(\log d)$, where d is the distance between x and f in the set (i.e, the difference between $\texttt{rank}(x)$ and $\texttt{rank}(f)$ in the set). This operation returns a finger to x if x is contained in the set, otherwise a finger to the largest element smaller than x in the set. This implementation supports the insertion of an element x immediately to the left or to the right of a finger in worst-case constant time.

In the description of the RAM Deferred Data Structure from Theorem 3, we substitute the dynamic predecessor and successor structure $\mathcal{B}$ by a finger search tree $\mathcal{F}_{\texttt{select}}$, as described by Brodal [7]. Once a block g is identified, every element in g is a valid pivot for the rest of the elements in $\mathcal{M}$. In order to capture this idea, we modify the structure $\mathcal{F}_{\texttt{select}}$ so that it contains blocks (i.e., a sequence of consecutive values) instead of singleton pivots. Each element in $\mathcal{F}_{\texttt{select}}$ points in $\mathcal{M}$ to the beginning and the end of the block g that it represents and in each run to the position where the elements of g partition the run. This modification allows the structure to answer `select` queries, taking advantage of the structure and order in the queries and of the structure and order of the input data. But in order to answer `rank` queries taking advantage of the features in the queries and the input data, the structure needs another finger search tree $\mathcal{F}_{\texttt{rank}}$. In $\mathcal{F}_{\texttt{rank}}$ the structure stores for each block g identified, the value of one of the elements in g, and pointers in $\mathcal{M}$ to the beginning and the end of g and in each run to the position where the elements of g partition the run. We name this structure Full-Synergistic Deferred Data Structure.

Theorem 4. *Consider a multiset $\mathcal{M}$ of size n formed by ρ runs and δ blocks. The* Full-Synergistic Deferred Data Structure *identifies γ blocks in the process of answering q online* **rank** *and* **select** *queries over $\mathcal{M}$. The q queries correspond to elements of* **ranks** *$r_1, \ldots, r_q$. Let $s_1, \ldots, s_\beta$ be the sizes of the β selection blocks determined by the γ blocks in all runs. Let $m_1, \ldots, m_\lambda$ be the numbers of pivot blocks among this selection blocks corresponding to the values of the λ pivots with multiplicity greater than 1, respectively. Let $\rho_0, \ldots, \rho_\xi$ be the sequence where ρ_i is the number of runs that have elements with values between the pivots i and $i+1$ sorted by* **ranks**, *for $i \in [1..\xi]$. Let $d_1, \ldots, d_{q-1}$ be the sequence where d_j is the number of identified blocks between the block that answers the $j-1$-th query and the one that answers the j-th query before starting*

the steps to answer the j-th query, for $j \in [2..q]$. Let u and $g_1, \ldots, g_u$ be the number of **rank** *queries and the sizes of the identified and searched blocks in the process of answering the u* **rank** *queries, respectively. The* Full-Synergistic Deferred Data Structure *answers the q online* **rank** *and* **select** *queries performing within $O(n+\sum_{i=1}^{\beta} \log s_i + \beta \log \rho - \sum_{i=1}^{\lambda} m_i \log m_i - \sum_{i=0}^{\xi} \rho_i \log \rho_i + \sum_{i=1}^{q-1} \log d_i + \sum_{i=1}^{u} \log g_i) \subseteq O\left(n \log n - \sum_{i=0}^{q} \Delta_i \log \Delta_i + q \log n\right)$ data comparisons, where $\Delta_i = r_{i+1} - r_i$, $r_0 = 0$ and $r_{q+1} = n$.*

Proof. The steps for answering a new $\texttt{select}(i)$ query are the same as the above description except when the algorithm searches for the nearest pivots to the left and right of the query position i. In this case, the algorithm searches for the position i in $\mathcal{F}_{\texttt{select}}$. If i is contained in an element of $\mathcal{F}_{\texttt{select}}$, then the block g that contains the element in the position i has already been identified. If i is not contained in an element of $\mathcal{F}_{\texttt{select}}$, then the returned finger f points the nearest block b to the left of i. The block that follows f in $\mathcal{F}_{\texttt{select}}$ is the nearest block to the right of i. Given f, the algorithm inserts in $\mathcal{F}_{\texttt{select}}$ each block identified in the process of answering the query in constant time and stores the respective pointers to positions in $\mathcal{M}$. In $\mathcal{F}_{\texttt{rank}}$ the algorithm searches for the value of one of the elements in g or b. Once the algorithm obtains the finger returned by this search, the algorithm inserts in $\mathcal{F}_{\texttt{rank}}$ the value of one of the elements of each block identified in constant time and stores the respective pointers to positions in $\mathcal{M}$.

The algorithm answers a new $\texttt{rank}(x)$ query by finding the *selection block* s_j in the j-th run such that x is between the smallest and the greatest value of s_j for all $j \in [1..\rho]$, similar to the steps of the RAM Deferred Data Structure for answering the query . For that the algorithm searches for the value x in $\mathcal{F}_{\texttt{rank}}$. The number of data comparisons performed by this searching process is within $O(\log d)$, where d is the number of blocks in $\mathcal{F}_{\texttt{rank}}$ between the last inserted or searched block and returned finger f. Given the finger f, there are three possibilities for the `rank` r of x: (i) r is between the `ranks` of the elements at the beginning and the end of the block pointed by f, (ii) r is between the `ranks` of the elements at the beginning and the end of the block pointed by the finger following f, or (iii) r is between the `ranks` of the elements in the selection blocks determined by f and the finger following f. In the cases (i) and (ii), a binary search inside the block yields the answer of the query. In case (iii), the algorithm applies the same steps as the algorithm `Quick Synergy MultiSelection` in order to compute the median μ of the middles and partitions the selection blocks by μ. The algorithm then decides to which side x belongs. These doubling searches identify two new blocks in the structure, the block that contains the greatest element smaller than or equal to x in $\mathcal{M}$ and the block that contains the smallest element greater than x in $\mathcal{M}$. Once compute $\texttt{rank}(x)$, the algorithm searches for this value in $\mathcal{F}_{\texttt{select}}$. It inserts then in $\mathcal{F}_{\texttt{select}}$ the block that contains the greatest element smaller than or equal to x and the block that contains the smallest element greater than x. □

The process of detecting the ϕ pivot positions seen in Section 2.5 allows the Full-Synergistic Deferred Data Structure to insert these pivots in $\mathcal{F}_{\texttt{select}}$ and $\mathcal{F}_{\texttt{rank}}$. For each pivot position p in $\mathcal{F}_{\texttt{select}}$ and $\mathcal{F}_{\texttt{rank}}$, the structure stores pointers to the end of the runs detected on the left of p; to the beginning of the runs detected on the right of p; and to the position of p in the multiset.

Corollary 5. *Let $\mathcal{M}$ be a multiset of size n with ϕ pivot positions. The ϕ pivot positions divide $\mathcal{M}$ into $\phi + 1$ sub-instances of size $n_0, \ldots, n_\phi$ (such that $\sum_{i=0}^{\phi} n_i = n$). Let q be the number of online* **rank** *and* **select** *queries over $\mathcal{M}$, such that q_i queries correspond to the sub-instance I_i, for $i \in [0..\phi]$. In each sub-instance I_i of size n_i formed by ρ_i runs, the* Full-Synergistic Deferred Data Structure *identifies γ_i blocks in the process of answering q_i online* **rank** *and* **select** *queries over I_i. Let $s_{i1}, s_{i2}, \ldots, s_{i\beta_i}$ be the sizes of the β_i selection blocks determined by the γ_i blocks in all runs of I_i. Let $m_{1i}, \ldots, m_{i\lambda_i}$ be the numbers of pivot blocks among this selection blocks corresponding to the values of the λ_i pivots with multiplicity greater than 1, respectively. Let $\rho_{i0}, \ldots, \rho_{i\xi_i}$ be the sequence where ρ_{ij} is the number of runs that have elements with values between the pivots ij and $i(j+1)$ sorted by* **ranks**, *for $j \in [1..\xi_i]$. Let $d_{i1}, d_{i2}, \ldots, d_{iq_{i-1}}$ be the sequence where d_{ij} is the number of identified blocks between the block that answers the $ij-1$-th query and the one that answers the ij-th query before starting the steps for answering the ij-th query, for $j \in [2..q_i]$. Let u_i and $g_{i1}, \ldots, g_{iu_i}$ be the number of* **rank** *queries and the sizes of the identified and searched blocks in the process of answering the u_i* **rank** *queries over I_i, respectively. There exists an algorithm that answers the q online* **rank** *and* **select** *queries performing within*

$O(n+\sum_{i=0}^{\phi}\left\{\sum_{j=1}^{\beta_i}\log s_{ij}+\beta_i\log\rho_i-\sum_{j=1}^{\lambda_i}m_{ij}\log m_{ij}-\sum_{j=0}^{\xi_i}\rho_{ij}\log\rho_{ij}+\sum_{j=1}^{q_i-1}\log d_{ij}+\sum_{j=1}^{u_i}\log g_{ij}\right\})$ *data comparisons.*

The Full-Synergistic Deferred Data Structure has two advantages over the RAM Deferred Data Structure: (i) it is in the pointer machine model of computation, which is less powerful than the RAM model; and (ii) it takes advantage of the structure and order in the queries and of the structure and order (local and global) in the input, when the RAM Deferred Data Structure does not take advantage of the order in the queries. Next, we present two compressed data structures, taking advantage of the block representation of a multiset $\mathcal{M}$ while supporting the operators `rank` and `select` on $\mathcal{M}$.

5 Compressed Data Structures for `Rank` and `Select`

We describe two compressed representations of a multiset $\mathcal{M}$ of size n formed by ρ runs and δ blocks while supporting the operators `rank` and `select` on it. The first compressed data structure represents $\mathcal{M}$ in $\delta\log\rho+3n+o(\delta\log\rho+n)$ bits and supports each `rank` query in constant time and each `select` query in time within $O(\log\log\rho)$. The second compressed data structure represents $\mathcal{M}$ in $\delta\log\delta+2n+O(\delta\log\log\delta)+o(n)$ bits and supports each `select` query in constant time and each `rank` query in time within $O\left(\frac{\log\delta}{\log\log\delta}\right)$.

Given a bitvector $\mathcal{V}$, $\texttt{rank}_1(\mathcal{V}, j)$ finds the number of occurrences of bit 1 in $\mathcal{V}[0..j]$, and $\texttt{select}_1(\mathcal{V}, i)$ finds the position of the i-th occurrence of bit 1 in $\mathcal{V}$. Given a sequence $\mathcal{S}$ from an alphabet of size ρ, $\texttt{rank}(\mathcal{S}, c, j)$ finds the number of occurrences of character c in $\mathcal{S}[0..j]$; $\texttt{select}(\mathcal{S}, c, i)$ finds the position of the i-th occurrence of character c in $\mathcal{S}$; and $\texttt{access}(\mathcal{S}, j)$ returns the character at position j in $\mathcal{S}$.

5.1 Rank-aware Compressed Data Structure for `Rank` and `Select`

The Rank-aware Compressed Data Structure supports `rank` in constant time and `select` in time within $O(\log\log\rho)$ using $\delta\log\rho+3n+o(\delta\log\rho+n)$ bits. It contains three structures $\mathcal{A}$, $\mathcal{B}$ and $\mathcal{C}$ representing bitvectors of size n supporting for $\mathcal{V}\in\{\mathcal{A},\mathcal{B},\mathcal{C}\}$, $\texttt{rank}_1(\mathcal{V}, j)$ and $\texttt{select}_1(\mathcal{V}, i)$ in constant time using $n+o(n)$ bits each [8]. It contains a data structure $\mathcal{S}$ representing a sequence of length δ from an alphabet of size ρ supporting $\texttt{rank}(\mathcal{S}, c, j)$ in time within $O(\log\log\rho)$, $\texttt{access}(\mathcal{S}, j)$ in time within $O(\log\log\rho)$, and $\texttt{select}(\mathcal{S}, c, i)$ in constant time, using $\delta\log\rho+o(\delta\log\rho)$ bits [14]. Given the blocks $g_1,\ldots,g_\delta$ in sorted order, $\mathcal{A}$ contains the information of the lengths of these blocks in this order in a bitvector of length n with a 1 marking the position where each block starts. $\mathcal{B}$ contains the information of the lengths of the blocks similar to $\mathcal{A}$ but with the blocks maintaining the original order, such that all blocks belonging to the same run are consecutive. $\mathcal{C}$ contains the information of the length of the runs in a bitvector of length n with a 1 marking the position where each run starts. The structure $\mathcal{S}$ contains the run to which g belongs for each block g in sorted order.

Theorem 5. *Let $\mathcal{M}$ be a multiset formed by ρ runs and δ blocks. The* Rank-aware Compressed Data Structure *represents $\mathcal{M}$ in $\delta\log\rho+3n+o(\delta\log\rho+n)$ bits supporting each* `rank` *query in constant time and each* `select` *query in time within $O(\log\log\rho)$.*

Proof. To answer a query $\texttt{rank}(\mathcal{M}, x)$, the following operations are executed: $\texttt{rank}_1(\mathcal{C}, i)$ returns the run r that contains x in constant time, where i is the position of x in the original order of $\mathcal{M}$; $\texttt{select}_1(\mathcal{C}, r)$ returns the position q where r starts in the original order of $\mathcal{M}$ in constant time; $\texttt{rank}_1(\mathcal{B}, i)-\texttt{rank}_1(\mathcal{B}, q-1)$ returns the position p inside of r of the block g that contains x in constant time; $\texttt{select}(\mathcal{S}, r, p)$ returns the position j of g in sorted order in constant time; and $\texttt{select}_1(\mathcal{A}, j)$ returns the `rank` of the first element in g in constant time.

For answering a query $\texttt{select}(\mathcal{M}, i)$, the following operations are executed: $\texttt{rank}_1(\mathcal{A}, i)$ returns the position j of the block g in sorted order that contains the selected element in constant time; $\texttt{access}(\mathcal{S}, j)$ returns the *run* r that contains the selected element in time within $O(\log\log\rho)$; $\texttt{rank}(\mathcal{S}, r, j)$ returns the position p of g inside r in time within $O(\log\log\rho)$; and $\texttt{select}_1(\mathcal{B}, p+\texttt{rank}_1(\mathcal{B}, \texttt{select}_1(\mathcal{C}, r)))$ returns the position where g starts in the original order of $\mathcal{M}$ in constant time. □

We describe next a compressed data structure that represents a multiset, taking advantage of the block representation of it, but unlike Rank-aware Compressed Data Structure, the structure supports `select` in constant time.

5.2 Select-aware Compressed Data Structure for `Rank` and `Select`

The Select-aware Compressed Data Structure supports `select` in constant time and `rank` in time within $O\left(\frac{\log \delta}{\log \log \delta}\right)$ using $\delta \log \delta + 2n + O(\delta \log \log \delta) + o(n)$ bits. It contains the same two structures $\mathcal{A}$ and $\mathcal{B}$ described above and a structure representing a permutation π of the numbers $[1..\delta]$ supporting the direct operator $\pi()$ in constant time and the inverse operator $\pi^{-1}()$ in time within $O\left(\frac{\log \delta}{\log \log \delta}\right)$ using $\delta \log \delta + O(\delta \log \log \delta)$ bits [24]. Given the blocks $g_1, \ldots, g_\delta$ in sorted order, $\pi(i)$ returns the position j of the block g_i in the original order of $\mathcal{M}$ and $\pi^{-1}(j) = i$ if the position of the block g_i is j in the original order of $\mathcal{M}$.

Theorem 6. *Let $\mathcal{M}$ be a multiset formed by ρ runs and δ blocks. The* Select-aware Compressed Data Structure *represents $\mathcal{M}$ in $\delta \log \delta + 2n + O(\delta \log \log \delta) + o(n)$ bits supporting each `select` query in constant time and each `rank` query in time within $O\left(\frac{\log \delta}{\log \log \delta}\right)$.*

Proof. To answer a query $\texttt{select}(\mathcal{M}, i)$, the following operations are executed: $\texttt{rank}_1(\mathcal{A}, i)$ returns the position j of the block g_j in sorted order that contains the selected element in constant time; $\pi(j)$ returns the position p of g_j in the original order of $\mathcal{M}$ in constant time; and $\texttt{select}_1(\mathcal{B}, p)$ returns the position where g_j starts in $\mathcal{M}$ in constant time.

To answer a query $\texttt{rank}(\mathcal{M}, x)$, the following operations are executed: $\texttt{rank}_1(\mathcal{B}, i)$ returns the position j of the block g that contains x in constant time, where i is the position of x in the original order of $\mathcal{M}$; $\pi^{-1}(j)$ returns the position p of g in sorted order in time within $O\left(\frac{\log \delta}{\log \log \delta}\right)$; and $\texttt{select}_1(\mathcal{A}, p)$ returns the `rank` of the first element of g in constant time. □

This concludes the description of our synergistic results. In the next section we discuss how these results relate to various past results and future work.

6 Discussion

In the context of deferred data structure, the concept of runs was introduced previously in [2, 19], but for a different purpose than the refined analysis of the complexity presented in this work: we clarify the difference and the research perspectives that it suggests in Section 6.1, and other perspectives for future research in Section 6.2. At a higher cognition level, we discuss the importance of categorizing techniques of multivariate analysis of algorithms in Section 6.3.

6.1 Comparison with previous work

Kaligosi et al.'s multiselection algorithm [19] and Barbay et al's deferred data structure [2] use the very same concept of *runs* as the one described in this work. The difference is that whereas we describe algorithms which *detect* the existing runs in the input in order to take advantage of them, the algorithms described by those previous works do not take into consideration any pre-existing runs in the input (assuming that there are none) and rather build and maintain such runs as a strategy to minimize the number of comparisons performed while partially sorting the multiset. We leave the combination of both approaches as a topic for future work which could probably shave a constant factor off the number of comparisons performed by the Sorting and MultiSelection algorithms and by the Deferred Data Structures supporting `rank` and `select` on Multisets. Johnson and Frederickson [13] described an algorithm answering a single `select` query in a set of sorted arrays of sizes $r_1, r_2, \ldots, r_\rho$, in time within $O(\sum_{i=1}^{\rho} \log r_i)$. Using their algorithm on pre-existing runs outperforms the Deferred Data Structures described in Section 4, when there is a single

query. Yet it is not clear how to generalize their algorithm into a deferred data structure in order to support more than one query. The difference is somehow negligible as the cost of such a query is anyway dominated by the cost ($n-1$ comparisons) of partitioning the input into runs. We leave the generalization of Johnson and Frederickson's algorithm into a deferred data structure which optimally supports more than one query as an open problem. We describe additional perspectives for future research in the next section.

6.2 Perspectives for future research

One question to tackle is to see how frequent are "easy" instances in concrete applications, in terms of input order and structure, and in terms of query order and structure; and how much advantage can be taken of them.

Barbay and Navarro [3] described how sorting algorithms in the comparison model directly imply encodings for permutations, and in particular how sorting algorithms taking advantage of specificities of the input imply compressed encodings of such permutations. By using the similarity of the execution tree of the algorithm `MergeSort` with the well known `Wavelet Tree` data structure, they described a compressed data structure for permutations taking advantage of local order, i.e., using space proportional to $\mathcal{H}(r_1, \ldots, r_\rho)$ and supporting `direct access` (i.e. $\pi()$) and `inverse access` (i.e. $\pi^{-1}()$) in worst time within $O(1 + \lg \rho)$ and average time within $O(1 + \mathcal{H}(r_1, \ldots, r_\rho))$. We leave the definition of a compressed data structure for multisets taking additional advantage of its structure and global order as future work.

Another perspective is to generalize the synergistic results to related problems in computational geometry: Karp et al. [20] defined the first deferred data structure not only to support `rank` and `select` queries on multisets, but also to support online queries in a deferred way on Convex Hull in two dimensions and online Maxima queries on sets of multi-dimensional vectors. One could refine the results from Karp et al. [20], expressed in function of the number of queries, to take into account the blocks between each queries (i.e., the structure in the queries) as Barbay et al. [2] did for multisets; but also for the relative position of the points (i.e., the structure in the data) as Afshani et al. [1] did for Convex Hulls and Maxima; the order in the points (i.e., the order in the data), as computing the convex hull in two dimension takes linear time if the points are sorted; and potentially the order in the queries.

6.3 Importance of the Parameterization of Structure and Order

The computational complexity of most problems is studied in the worst case over instances of fixed size n, for n asymptotically tending to infinity. This approach was refined for NP-difficult problems under the term "parameterized complexity" [12], for polynomial problems under the term "Adaptive Algorithms" [11, 23], and more simply for data encodings under the term of "Data Compression" [3], for a wide range of problems and data types. Such a variety of results has motivated various tentative to classify them, in the context of NP-hard problems with a theory of Fixed Parameter Tractability [12], and in the context of sorting in the comparison model with a theory of reduction between parameters [26]. We introduced two other perspectives from which to classify algorithms and data structures. Through the study of the sorting of multisets according to the potential "easiness" in both the order and the values in the multiset, we aimed to introduce a way to classify refined techniques of complexity analysis between the ones considering the input order and the ones considering the structure in the input; and to show an example of the difficulty of combining both into a single hybrid algorithmic technique. Through the study of the online support of `rank` and `select` queries on multisets according to the potential "easiness" in both the order and the values in the queries themselves (in addition to the potential easiness in the data being queried), we aimed to introduce such categorizations. We predict that such analysis techniques will take on more importance in the future, along with the growth of the block between practical cases and the worst case over instances of fixed sizes. Furthermore, we conjecture that synergistic techniques taking advantage of more than one "easiness" aspect will be of practical importance if the block between theoretical analysis and practice is to ever be reduced.

References

1. Afshani, P., Barbay, J., Chan, T.M.: Instance-optimal geometric algorithms. In: Proceedings of the Annual IEEE Symposium on Foundations of Computer Science (FOCS). pp. 129–138. IEEE Computer Society (2009)
2. Barbay, J., Gupta, A., Satti, S.R., Sorenson, J.: Near-optimal online multiselection in internal and external memory. Journal of Discrete Algorithms (JDA) 36, 3–17 (2016)
3. Barbay, J., Navarro, G.: On compressing permutations and adaptive sorting. Theoretical Computer Science (TCS) 513, 109–123 (2013)
4. Beame, P., Fich, F.E.: Optimal bounds for the predecessor problem and related problems. Journal of Computer and System Sciences (JCSS) 65(1), 38 – 72 (2002)
5. Bentley, J.L., Yao, A.C.C.: An almost optimal algorithm for unbounded searching. Information Processing Letters (IPL) 5(3), 82–87 (1976)
6. Blum, M., Floyd, R.W., Pratt, V.R., Rivest, R.L., Tarjan, R.E.: Time bounds for selection. Journal of Computational System Science (JCSS) 7(4), 448–461 (1973)
7. Brodal, G.S.: Finger search trees with constant insertion time. In: Proceedings of the ninth annual ACM-SIAM symposium on Discrete algorithms (SODA). pp. 540–549. Society for Industrial and Applied Mathematics (1998)
8. Clark, D.R.: Compact Pat Trees. Ph.D. thesis, University of Waterloo (1996)
9. Demaine, E.D., López-Ortiz, A., Munro, J.I.: Adaptive set intersections, unions, and differences. In: Proceedings of the 11^{th} ACM-SIAM Symposium on Discrete Algorithms (SODA). pp. 743–752 (2000)
10. Dobkin, D.P., Munro, J.I.: Optimal time minimal space selection algorithms. Journal of the ACM (JACM) 28(3), 454–461 (1981)
11. Estivill-Castro, V., Wood, D.: A survey of adaptive sorting algorithms. ACM Computing Surveys (ACMCS) 24(4), 441–476 (1992)
12. Flum, J., Grohe, M.: Parameterized Complexity Theory (Texts in Theoretical Computer Science. An EATCS Series). Springer-Verlag New York, Inc., Secaucus, NJ, USA (2006)
13. Frederickson, G.N., Johnson, D.B.: Generalized selection and ranking. In: Proceedings of the 12th Annual ACM Symposium on Theory of Computing (STOC), April 28-30, 1980, Los Angeles, California, USA. pp. 420–428 (1980)
14. Golynski, A., Munro, J.I., Rao, S.S.: Rank/select operations on large alphabets: A tool for text indexing. In: Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithm (SODA). pp. 368–373. SODA '06, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA (2006)
15. Guibas, L.J., McCreight, E.M., Plass, M.F., Roberts, J.R.: A new representation for linear lists. In: Proceedings of the ninth annual ACM symposium on Theory of computing (STOC). pp. 49–60. ACM Press, New York, NY, USA (1977)
16. Hoare, C.A.R.: Algorithm 64: Quicksort. Communication of the ACM (CACM) 4(7), 321 (1961)
17. Hoare, C.A.R.: Algorithm 65: Find. Communication of the ACM (CACM) 4(7), 321–322 (1961)
18. Huffman, D.A.: A method for the construction of minimum-redundancy codes. Proceedings of the Institute of Radio Engineers (IRE) 40(9), 1098–1101 (September 1952)
19. Kaligosi, K., Mehlhorn, K., Munro, J.I., Sanders, P.: Towards optimal multiple selection. In: Proceedings of the International Conference on Automata, Languages, and Programming (ICALP). pp. 103–114 (2005)
20. Karp, R.M., Motwani, R., Raghavan, P.: Deferred data structuring. SIAM Journal on Computing (SICOMP) 17(5), 883–902 (1988)
21. Knuth, D.E.: The Art of Computer Programming, Vol 3, chap. Sorting and Searching, Section 5.3. Addison-Wesley (1973)
22. Mannila, H.: Measures of presortedness and optimal sorting algorithms. IEEE Trans. Computers 34(4), 318–325 (1985)
23. Moffat, A., Petersson, O.: An overview of adaptive sorting. Australian Computer Journal (ACJ) 24(2), 70–77 (1992)
24. Munro, J.I., Raman, R., Raman, V., S., S.R.: Succinct representations of permutations and functions. Theoretical Computer Science (TCS) 438, 74 – 88 (2012)
25. Munro, J.I., Spira, P.M.: Sorting and searching in multisets. SIAM Journal on Computing (SICOMP) 5(1), 1–8 (1976)
26. Petersson, O., Moffat, A.: A framework for adaptive sorting. Discrete Applied Mathematics (DAM) 59, 153–179 (1995)
27. Takaoka, T.: Minimal mergesort. Tech. rep., University of Canterbury (1997), http://ir.canterbury.ac.nz/handle/10092/9676, last accessed [2016-08-23 Tue]

28. Takaoka, T.: Partial solution and entropy. In: Královič, R., Niwiński, D. (eds.) Mathematical Foundations of Computer Science (MFCS) 2009: 34th International Symposium, Novy Smokovec, High Tatras, Slovakia, August 24-28, 2009. Proceedings. pp. 700–711. Springer Berlin Heidelberg, Berlin, Heidelberg (2009)